

CEMETR-2021-06
JUNE 2021

CEME

Technical Report

The Center for Educational Measurement and Evaluation

Examining Inter-Rater Reliability of Evaluators
Judging Teacher Performance:
Proposing an Alternative to Cohen's Kappa

Richard G. Lambert

T. Scott Holcomb

Bryndle L. Bottoms

A PUBLICATION OF
THE CENTER FOR
EDUCATIONAL
MEASUREMENT
AND EVALUATION

**Examining Inter-Rater Reliability of Evaluators Judging Teacher Performance:
Proposing an Alternative to Cohen's Kappa**

Richard Lambert

Scott Holcomb

Bryndle Bottoms

Center for Educational Measurement and Evaluation

UNC Charlotte

Paper presented to the virtual Annual Meeting of the National Council on Measurement in Education, June, 2021.

Abstract

The validity of the Kappa coefficient of chance-corrected agreement has been questioned when the prevalence of specific rating scale categories is low and agreement between raters is high. The researchers proposed the Lambda Coefficient of Rater-Mediated Agreement as an alternative to Kappa to address these concerns. Lambda corrects for chance agreement based on specific assumptions about raters and the rater-mediated assessment process including rater-specific tendencies for strict or lenient ratings. Actual ratings of teacher profiles from an inter-rater reliability exercise confirmed the shortcomings of Kappa. The rater data also demonstrated the robustness of Lambda-1, Lambda-2, Gwet's AC1, and Gwet's AC2 to the data conditions that are problematic for Kappa. All four alternative chance-corrected agreement coefficients showed less variability across the 45 raters than Kappa. However, AC-2 was undetermined for 39 of the 45 raters. Simulation data demonstrated the robustness of the Lambda Coefficient of Rater-Mediated Agreement to the data conditions that are problematic for Kappa.

Examining Inter-Rater Reliability of Evaluators Judging Teacher Performance: An Alternative to Cohen's Kappa

Cohen's Kappa (Cohen, 1960) and Weighted Kappa (Cohen, 1968) are widely used measures of chance-corrected agreement between raters. Various questions have been raised about whether Kappa is actually correcting for chance agreement, whether it is useful for identifying and separating various sources of disagreement, the validity of Kappa coefficients when prevalence of specific categories on a rating scale is low, the validity of Kappa coefficients when agreement is high, and the generalizability of Kappa coefficients across populations and study conditions (Thompson & Walter, 1988; Feinstein & Cicchetti; 1990; Cicchetti & Feinstein, 1990). Gwet (2008) introduced AC1 and AC2 as alternatives to Kappa and demonstrated they are robust indexes not susceptible to the identified shortcomings of Kappa.

This study proposes an alternative to Kappa, Weighted Kappa, AC1, and AC2 that is rooted in theory regarding rater-mediated assessment (Engelhard & Wind, 2018). The Lambda Coefficient of Rater-Mediated Agreement is designed for use with ordinal scales that are often used to evaluate teacher performance. It examines inter-rater agreement corrected for the probability that raters may agree with expert raters by chance due to the response process they employ when they are uncertain about how to place a teacher on a rubric.

Rater-Mediated Assessment Theory

Most research on rater cognition focuses on the mental processes used by raters of student or examinee performances. The rating process followed by the raters judging constructed responses is intricate, however the judgment of teacher performance can pose even greater complexities. Similar to other examples of rater-mediated assessments, an observers' level of expertise, and the overall scoring task demands, drive ratings of teachers (Bell et al., 2018; Suto,

2012). Understanding rater cognition is crucial to making a validity argument to support the use of any rater-mediated assessment measure. According to Standard 1.12, which addresses “evidence regarding cognitive processes”, in the *Standards for Educational and Psychological Testing* (2014):

“If the rationale for score interpretation for a given use depends on premises about the psychological processes or cognitive operations of test takers, then theoretical or empirical evidence in support of those premises should be provided. When statements about the processes employed by observers or scorers are part of the argument for validity, similar information should be provided.” (p. 26)

It is common to provide raters of any construct a rubric and a set of exemplar and non-exemplar responses to serve as anchors or benchmarks in scoring decisions. This standard makes it important to go beyond rater training to analyze and interpret raters’ cognitive processes, and to make sure those processes are congruent with the measured construct (Bejar, 2012).

Engelhard et al. (2018) asserted the validity, reliability, and fairness of rater-mediated assessments relies on both the quality of the rater’s cognitive process and psychometric properties of the measure. In rater-mediated assessment, understanding raters’ scoring processes is an important component in understanding what is actually being measured by an assessment (Crisp, 2012). Interpreting and anticipating a rater’s cognitive process “can provide practical information to assist those who are designing performance tasks and rubrics, selecting raters, training raters, and developing quality control procedures to monitor rater performance, particularly in ‘real time’ as a scoring session is proceeding” (Myford, 2012, p. 49).

The lens model proposed by Brunswik (1952) was initially designed as a human judgment and decision-making conceptual framework. This model was adapted by Engelhard

(2013) as a conceptual framework for rater judgment and decision-making. The goal of this adapted version of the lens model “is to have a close correspondence between the latent variable and the observed ratings” (Engelhard & Wind, 2018, p. 81). The connection takes place between the items, rater, and rating scale. Teacher evaluation is multidimensional and requires raters to have a high level of expertise and skill in interpreting information through this model.

The following set of assumptions about raters, and the complex response process they use to arrive at ratings, serve as a theoretical foundation for the Lambda Coefficient of Rater-Mediated Agreement. We posit the following principles regarding the internal cognitive process raters employ when they are confident about which rating to assign:

- Raters are trained evaluators and function as expert professionals.
- Rather than acting as scoring machines, raters bring their own experiences and expertise to the rating process.
- Raters use a complex, three-stage internal response process to make ratings.
- First, raters acquire an overall impression, based on global evidence, to arrive at a starting point on a rubric or rating scale.
- Second, raters synthesize information from previous ratings, analyze observational data, and interpret evidence and artifacts.
- Third, raters combine their overall impressions with their analysis of evidence to settle on a final placement on a rubric or rating scale.
- A rater’s individual tendencies toward strictness and leniency influence this complex internal response process.

This process functions at several levels. When raters consider how to make a rating on an item that addresses a specific area of practice, such as ratings focused on particular

competencies, they start with their overall impression, analyze a variety of item-specific pieces of evidence, synthesize the ratings they have made across items that address similar content, and then settle on a final rating. Similarly, when raters make global ratings, such as ratings of overall effectiveness or quality, they may start with their overall impression, analyze a variety of pieces of evidence, synthesize the ratings they have made across items that address various content areas, and then settle on a final rating. However, raters do resort to guessing, or at least a random process similar to guessing, when they are uncertain about a particular rating. We posit the following principles regarding the internal cognitive process raters employ when they are uncertain about which rating to assign:

- Professional raters can be, on occasion, uncertain about their selection of ratings.
- A professional rater may, on occasion, lack the experience, expertise, or evidence to have confidence in a particular rating.
- When uncertain, raters make ratings by a random process that mimics the three-stage internal cognitive response process they use when confident in their ratings. These ratings may, by chance, agree with the ratings of another rater or those from an expert panel.
- When uncertain, raters select a random starting point for deliberations.
- When uncertain, raters may synthesize previous ratings and analyze evidence, but this process does not resolve their uncertainty. When uncertain, raters may have little confidence in their previous ratings and may lack sufficient evidence to support a particular rating.
- When uncertain, raters combine their initial random starting point with their inconclusive analysis of evidence to settle on a final rating.

- A rater's individual tendencies toward strictness and leniency influence this random response process.

We developed the Lambda Coefficient of Rater-Mediated Agreement based on these assumptions concerning the response process raters use when applying ordinal ratings scales to tasks such as teacher evaluation.

The Lambda Coefficient of Rater-Mediated Agreement

Cohen (1960) introduced the Kappa coefficient of chance-corrected agreement. Kappa is equivalent to the proportion of the ratings that are in agreement with another rater, after removing the proportion of the agreement ratings that may have occurred by chance. The formulae take the following forms:

$$\kappa = (p_a - p_e) / (1 - p_e) \quad (1)$$

$$\sigma_\kappa = \sqrt{ \{ [p_a(1 - p_a)] / [n(1 - p_e)^2] \} } \quad (2)$$

Where:

p_a = Proportion of exact agreement.

p_e = Expected proportion chance agreement. For Kappa, this quantity is equal to the sum of the products of the marginal proportions associated with each cell.

n = number of ratings.

There have been various alternatives to Kappa proposed in the years since (Holley & Guilford, 1964; Maxwell, 1970; Krippendorff, 1970; Jason & Vegelius, 1979; Brennan & Prediger, 1981; Perreault & Leigh, 1989; Bryt, Bishop, & Carlin, 1993; Gwet, 2008). In addition, Bennett, Albert, and Goldstein (1954) proposed a method equivalent to the method reintroduced by Bryt, Bishop, & Carlin (1993) prior to the introduction of Kappa. These alternatives to Kappa were developed based on the assumptions of generalizability theory, applications for nominal

scales, or both. Their focus was primarily on agreement among raters, judges, or observers with respect to the presence or absence of specific characteristics, symptoms, or diagnoses. Such applications do not involve rater strictness or leniency as is often present in rater use of ordinal rating scales.

The researchers are proposing the Lambda Coefficient of Rater-Mediated Agreement for a different set of applications. The researchers are proposing the Lambda based on the theoretical propositions of rater-mediated assessment (Engelhard & Wind, 2018) and our applied work with teacher performance evaluations. Teacher performance evaluations are typically conducted using ordinal scales. The agreement of interest is between individual raters and expert raters.

Furthermore, teacher performance evaluations are high stakes endeavors, and strict or lenient ratings can have significant consequences for teachers. Evaluators make placements on such scales based on observational data, classroom artifacts, student work samples, and general overall impressions. Raters can and will have their own personal tendencies toward strictness or leniency, or even biases. In addition, teacher evaluators can be uncertain about a particular rating and can use a random process to settle on their final placements on ordinal rating scales. Our goal is to correct for chance agreement that may occur due to this complex cognitive process.

We set out to develop a coefficient for ordinal scales that meets several important criteria. First, Lambda had to agree with Kappa when all ratings fall on the main diagonal of the ratings matrix. When all rater placements are in agreement with the expert ratings, Lambda and Kappa both = 1.0. Furthermore, when rater agreement, strictness, and leniency are all equal this is equivalent to a rater cognitive process that involves simple guessing. Therefore, Kappa and Lambda should agree in these circumstances and they do. Next, we sought a coefficient that would equal zero when all ratings in the ratings matrix have equal frequency, and both Kappa

and Lambda equal zero under these circumstances. Finally, we sought to develop a coefficient that has a reasonable upper bound on the magnitude of the correction for chance, similar to Gwet's approach (2008). Lambda-1, described below, met all of these criteria. For example, Lambda-1 has an upper bound on chance agreement of .5 for a 4x4 ratings matrix as defined by these quantities.

$$p_e \leq 2L / q \quad (3)$$

$$p_e \leq 2S / q \quad (4)$$

Where:

p_e = Expected proportion chance agreement.

L = proportion of ratings that are lenient, or above the "correct" or "expert" rating.

S = proportion of ratings that are strict, or below the "correct" or "expert" rating.

q = number of steps on the ordinal rating scale.

The general form for λ , applicable to both λ_1 and λ_2 , and to rating scales with any number of steps, can be expressed as:

$$\lambda = (p_a - p_e) / (1 - p_e) \quad (5)$$

$$p_e = \sum p_s p_c p_f \quad (6)$$

$$\sigma_\lambda = \sqrt{ \{ [p_a(1 - p_a)] / [n(1 - p_e)^2] \} } \quad (7)$$

Where:

p_a = Proportion of exact agreement.

p_e = Proportion expected chance agreement.

Σ = Sum across all cells from $r=1, c=1$ to $r=q, c=q$.

r = row.

c = column.

n = number of ratings.

q = The number of steps on the ordinal rating scale.

p_s = Probability of picking the given cell as a starting point (s) for deliberation.

p_c = Proportion of ratings for which the given column (c) is used as a correct answer.

p_f = Expected probability of exact agreement when the given cell was used as a starting point, and the rater makes a final (f) rating informed by their tendency for agreement, strictness, and leniency.

The only difference between λ_1 and λ_2 is the formula for p_s . For λ_1 , $p_s = 1/q$. This value assumes the rater is uncertain about which rating to give, arrives at a random starting point for their deliberations, and is equally likely to select any of the points on the rating scale as a starting point. For λ_2 , p_s is set to the proportion of the total ratings for which the rater used the rating scale level associated with the cell in question. This value is the marginal proportion for the given row in the matrix. This value also assumes the rater is uncertain about which rating to give, uses guessing as a means to arrive at a starting point for their deliberations, and their internal guessing process weights the points on the rating scale according to how frequently they use each point. So for example, if a rater very rarely uses a particular point on the rating scale, λ_2 assumes the rater would be much less likely to select that point as a starting point for deliberations. To illustrate how Lambda-1 and Lambda-2 work in practice, see Figure 1 for a two-point ordinal rating scale. See Figure 2 for a three-point ordinal rating scale and Figure 3 for a four-point ordinal rating scale. Just for illustration purposes, we have included category labels that might apply to a teacher performance evaluation rubric.

The Current Study

The purpose of this study was twofold. First, the researchers sought to test the performance of Lambda relative to Kappa, AC-1, and AC-2 using field data. Second, the researchers sought to evaluate Lambda relative to Kappa, AC-1, and AC-2 with simulated data that represents the high agreement / low frequency of specific categories data conditions under which Kappa is known to yield paradoxical results. Specifically, this study examined the following research questions:

1. How does the Lambda Coefficient of Rater-Mediated Agreement perform relative to Kappa, AC-1, and AC-2 given real world teacher performance evaluation data?
2. Does the Lambda Coefficient of Rater-Mediated Agreement yield chance-corrected coefficients of agreement that are robust to data conditions that have been shown to be problematic for Kappa (high agreement and some rating scale categories with low prevalence)?

Methods

Evaluators charged with conducting state-mandated performance evaluations of all licensed pre-kindergarten teachers working in non-public school settings within one state participated in an inter-rater reliability certification exercise. Evaluators ($n=45$) made placements on five progressions across each of 10 online teacher profiles for a total of 2,250 ratings. The evaluators rated the teacher profiles using the North Carolina Teacher Evaluation Process rubric, which is the same teacher performance evaluation measure used to conduct evaluations across the entire state. The measure includes five progressions, called standards, each of which contains specific behavioral anchors and is supported by a series of rubrics called “elements.” The evaluators used the same four-point rating scale for all standards and elements. The ordinal scale points were labeled (1) Developing, (2) Proficient, (3) Accomplished, and (4) Distinguished.

Agreement was evaluated against the “correct answer” ratings from a panel of five experts who had previously achieved consensus. The following statistics were calculated for each evaluator using the criteria for exact agreement with the expert panel: a.) agreement, leniency, and strictness percentages, b.) Cohen’s Kappa, c.) Lambda-1 Coefficient of Rater-Mediated Agreement, d.) Lambda-2 Coefficient of Rater-Mediated Agreement, e.) Gwet’s AC-1, and f.) Gwet’s AC-2.

An alternative scoring strategy that allowed for agreement between some adjacent ratings was also developed. Adjacent agreement was defined as an expert panel rating of “Proficient” and an evaluator rating of either “Proficient” or “Accomplished”, or an expert panel rating of “Accomplished” and an evaluator rating of either “Proficient” or “Accomplished.” Exact agreement was still required for expert panel ratings of either “Developing” or “Distinguished.” The rationale was there is no difference in how teachers who are rated as “Proficient” or “Accomplished” are treated within either a mentoring or a performance evaluation context in the particular state under investigation. Teachers must obtain ratings of at least “Proficient” across all standards by the end of their third year of teaching. Therefore, teachers rated as “Developing” receive additional mentoring. Teachers rated “Distinguished” are rare and may be asked to serve as model teachers, mentors, or evaluators. The same statistics were calculated for each evaluator using the criteria for adjacent agreement with the expert panel: a.) agreement, leniency, and strictness percentages, b.) Cohen’s Kappa, c.) Lambda-1 Coefficient of Rater-Mediated Agreement, d.) Lambda-2 Coefficient of Rater-Mediated Agreement, e.) Gwet’s AC-1, and f.) Gwet’s AC-2. It should be noted that the Kappa coefficient for the adjacent agreement condition is equivalent to a special case of Weighted Kappa (Cohen, 1968) with weights assigned according to this particular adjacent scoring scheme.

Results

First, we examined the distribution of the agreement, strictness, and leniency percentages for all 45 raters across both the exact and adjacent agreement conditions. Table 1 contains the mean, standard deviation, and five number summary for each of these percentages. The mean percent exact agreement across the 45 evaluators was 68.7% (SD=10.2) and values ranged from 42.0% to 88.0%. The mean percent lenient for exact agreement was 7.7% (SD=6.5) and values ranged from 0.0% to 26.0%. The mean percent strict for exact agreement was 23.6% (SD=12.0) and values ranged from 4.0% to 58.0%. Therefore, the raters as a group displayed moderate levels of agreement. However, a substantial minority of raters ($n = 7$) agreed with the expert panel for less than 60% of their ratings. The raters as a group also displayed more strictness in their ratings than they did leniency.

As expected, agreement percentages increased, and strictness and leniency percentages decreased, for the adjacent agreement condition. Only four of the 45 raters displayed adjacent agreement percentages less than 80%. The mean percent adjacent agreement across the 45 evaluators was 87.6% (SD=7.0) and values ranged from 64.0% to 100.0%. The mean percent lenient for adjacent agreement was 2.4% (SD=2.8) and values ranged from 0.0% to 10.0%. The mean percent strict for adjacent agreement was 10.0% (SD=7.4) and values ranged from 0.0% to 36.0%. For the adjacent agreement condition, strictness was again greater than leniency; however, both values were much lower than they were in the exact agreement condition.

The distributions of each of five coefficients of chance-corrected agreement were compared to address research question one. Table 2 contains the mean, standard deviation, and five number summary for each of the coefficients were calculated. The mean Kappa for exact agreement was .486 (SD=.160) and values ranged from .121 to .790. The mean Lambda-1

Coefficient of Rater-Mediated Agreement for exact agreement was .578 (SD=.140) and values ranged from .199 to .840. The mean Lambda-2 Coefficient of Rater-Mediated Agreement for exact agreement was .526 (SD=.137) and values ranged from .225 to .797. The mean Gwet's AC-1 for exact agreement was .606 (SD=.129) and values ranged from .273 to .852. The mean Gwet's AC-2 for exact agreement was .472 (SD=.092) and values ranged from .356 to .548. However, Gwet's AC-2 was undetermined for 39 of the 45 raters. Figure 4 displays these distributions as boxplots. AC-2 was not included due to small sample size ($n = 6$). The boxplots show that Lambda-1, Lambda-2, and Gwet's AC-1 yielded less severe and less variable corrections for chance agreement than Kappa. Lambda-1 showed sensitivity to one outlier rater not detected by the other coefficients.

A very consistent rank order of correction among the alternatives to Kappa emerged for the exact agreement condition (see Figure 6). Lambda-2 was closest to Kappa for 44 of the 45 raters. Lambda-1 yielded coefficients that were consistently higher than Kappa and Lambda-2 and lower than AC-1 and this pattern held for 42 of the 45 raters. AC-1 emerged as yielding the consistently highest coefficients and this pattern held for 43 of the 45 raters. AC-2 was not included in these comparisons due to the small sample size. As seen in Figure 6, the most dramatic changes to the relatively consistent rank order of the coefficients appear between Lambda-1 and AC-1 as indicated by several lines that break from the overall pattern.

The mean Kappa for adjacent agreement was .666 (SD=.143) and values ranged from .315 to 1.000 (see Table 3). The mean Lambda-1 Coefficient of Rater-Mediated Agreement for adjacent agreement was .812 (SD=.107) and values ranged from .442 to 1.000. The mean Lambda-2 Coefficient of Rater-Mediated Agreement for adjacent agreement was .714 (SD=.115) and values ranged from .430 to 1.000. The mean Gwet's AC-1 for adjacent agreement was .848

(SD=.090) and ranged from .532 to 1.000. The mean Gwet's AC-2 for adjacent agreement was .787 (SD=.052) and values ranged from .687 to .822. However, again Gwet's AC-2 was undetermined for 39 of the 45 raters. Figure 5 displays these distributions as boxplots. AC-2 was not included due to small sample size ($n = 6$). The boxplots show that Lambda-1, Lambda-2, and Gwet's AC-1 yielded less severe and less variable corrections for chance agreement than Kappa. Lambda-1 and AC-1 yielded similar distributions, displayed less variability than Kappa or Lambda-2, and were both sensitive to one outlier rater not detected by the other coefficients.

A very consistent rank order of correction among the alternatives to Kappa emerged for the adjacent agreement condition as well (see Figure 7). Lambda-2 was closest to Kappa for 44 of the 45 raters. Lambda-1 coefficients were higher than Kappa and Lambda-2 and lower than AC-1 for 44 of the 45 raters. AC-1 yielded the highest coefficients for 44 of the 45 raters. AC-2 was not included in these comparisons due to the small sample size ($n = 6$). There was only one exception to this pattern. One rater had 100% adjacent agreement with the expert panel and all coefficients equaled 1.00.

A simulation study addressed research question two. The simulation design extended the approach of Xie (2013) to include Lambda-1 and Lambda-2. For the purpose of this simulation study, we defined the Bias Index as Strictness minus Leniency (expressed as proportions). We defined the Prevalence Index as the proportion of rater selections using the lowest point on the ratings scale minus the proportion of rater selections using the highest point on the rating scale. We varied the Prevalence Index across all possible values for each condition. Four simulated conditions used a four point rating scale similar to the real world data conditions reported for research question one. These four conditions included high agreement and low category frequency conditions known to be problematic for Kappa. The four conditions were: 1.)

Agreement = 95%, Bias Index = .05, Prevalence ranged from .95 to -.95, 2.) Agreement = 90%, Bias Index = .10, Prevalence ranged from .90 to -.90, 3.) Agreement = 85%, Bias Index = .15, Prevalence ranged from .85 to -.85, and 4.) Agreement = 80%, Bias Index = .20, Prevalence ranged from .80 to -.80. We calculated Kappa, Lambda-1, and Lambda-2 for each of the four conditions across the applicable range of the Prevalence Index.

For condition 1 (95% agreement), the mean Kappa was .820 (SD = .147) and values range from -.053 to .904. The mean Lambda-1 was .933 (SD = .001) and values range from .932 to .934. The mean Lambda-2 was .842 (SD = .094) and values range from .487 to .907. Therefore, Lambda-1 yielded very consistent values (see Figure 8), Kappa yielded very inconsistent and over-corrected values, and Lambda-2 followed a similar pattern as Kappa but did not over-correct as much when the Prevalence Index was high. The remaining three conditions yielded similar patterns. For condition 2 (90% agreement), the mean Kappa was .739 (SD = .111) and values range from .298 to .825. The mean Lambda-1 was .867 (SD = .002) and values range from .863 to .871. The mean Lambda-2 was .752 (SD = .090) and values range from .474 to .826. For condition 3 (85% agreement), the mean Kappa was .674 (SD = .095) and values range from .355 to .756. The mean Lambda-1 was .801 (SD = .005) and values range from .793 to .810. The mean Lambda-2 was .683 (SD = .080) and values range from .460 to .755. For condition 4 (80% agreement), the mean Kappa was .621 (SD = .082) and values range from .370 to .696. The mean Lambda-1 was .737 (SD = .008) and values range from .722 to .750. The mean Lambda-2 was .625 (SD = .070) and values range from .444 to .691. For all four conditions, Lambda-2 tended to over-correct less than Kappa for high values of the Prevalence Index, converge with Kappa as the Prevalence Index got smaller, and yield nearly identical

values to Kappa at the minimum values of the Prevalence Index. Lambda-1 remained very consistent within each of the simulated conditions.

Discussion

The results of this study confirmed and extended previous research (Gwet, 2008) by illustrating the shortcomings of Kappa as a measure of chance-corrected agreement and the robustness of AC-1 to the data conditions associated with these shortcomings. These results also illustrated how the proposed Lambda-1 Coefficient of Rater-Mediated Agreement is resistant to the data conditions that are problematic for Kappa, and offers a slightly more conservative, less variable measure of chance-corrected agreement than AC-1 while also demonstrating greater sensitivity to outlier raters.

The data from this study contained various examples, both real and simulated, of the high agreement / low frequency of specific rating scale categories problem. For example, the real world data included very infrequent use of the “Distinguished” category by the evaluators or experts. In practice, raters use “Distinguished” very rarely and reserve its use for truly exceptional teachers. The paradoxical performance of Kappa found in previous studies under these data conditions was confirmed (Cicchetti and Feinstein, 1990; Gwet, 2008). Consistent with previous research, Cohen’s Kappa was overly sensitive and over-corrected when agreement was high and there was low frequency of specific categories on the rating scale. Cohen’s Kappa yielded results consistent with the paradox problem in which percent agreement is high and Cohen’s Kappa is low or even .000. However, both Lambda-1 Coefficient of Rater-Mediated Agreement and Gwet’s AC-1 yield results that were robust to both of these data conditions.

Separate examinations of the fictitious teacher profiles revealed several stark examples of this pattern. For example, neither the expert panel nor the evaluators selected ratings of

“Accomplished” or “Distinguished” for Profile 4. Across all 45 raters, agreement was high (93.33%) for both the exact and adjacent methods. However, Cohen’s Kappa was .000 for the exact method while Lambda-1 equaled .913 and AC-1 equaled .932. Cohen’s Kappa was also .000 for the adjacent method while Lambda-1 equaled .903 and AC-1 equaled .933. Similarly, for Profile 5 no ratings of “Developing” or “Distinguished” were selected by either the expert panel or the evaluators, and agreement was moderate (63.11%) for exact method and high for the adjacent method (96.44%). However, Cohen’s Kappa was .000 for the exact method while Lambda-1 equaled .508 and AC-1 equaled .589. Cohen’s Kappa was also .000 for the adjacent method while Lambda-1 equaled .947 and AC-1 equaled .964.

It is important to point out that Lambda is not meant to provide the rich information that a more complex measurement model can provide about individual raters and their tendencies. For example, the Many-Facets Rasch Model (Linacre, 1989) can provide a detailed calibration of individual rater strictness and leniency and potential biases. Lambda is a single coefficient and is agnostic to where in the rating space strictness or leniency occurs. It cannot detect or identify the steps on a rating scale that are associated with a rater’s tendencies for strictness or leniency. It is, however, useful as a red flag, as one indicator among many, of the need to support, retrain, or recertify individual raters.

In conclusion, this study confirmed the advantages of AC-1 over Kappa demonstrated in previous research (Gwet, 2008). In addition, this study introduced the Lambda Coefficient of Rater-Mediated Agreement. Lambda is rooted in the theoretical underpinnings of rater-mediated assessment (Engelhard & Wind, 2018). It operationalizes a series of proposed principles regarding the complex process by which raters make placements on ordinal progressions. Future research is needed to test these theoretical propositions and to investigate the cognitive processes

raters use when they feel confident in their ratings and when they are uncertain. The current study, with both field data and simulated data, highlighted the robustness of the Lambda Coefficient of Rater-Mediated Agreement to the data conditions that are problematic for Kappa. Future research is needed to test Lambda across a wider range of field and simulated data conditions.

References

- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education (Eds.). (2014). *Standards for educational and psychological testing*. American Educational Research Association.
- Bejar, I. I. (2012). Rater cognition: Implications for validity. *Educational Measurement: Issues and Practice*, 31(3), 2-9.
- Bell, C. A., Jones, N. D., Qi, Y., & Lewis, J. M. (2018). Strategies for assessing classroom teaching: Examining administrator thinking as validity evidence. *Educational Assessment*, 23(4), 229-249. <https://doi.org/10.1080/10627197.2018.1513788>
- Bennett E. M., Albert R., Goldstein A. C. (1954). Communications through limited response questioning. *Public Opinion Quarterly*, 18 (3), 303-308.
- Brennan, R. L., & Prediger, D. J. (1981). Coefficient Kappa: some uses, misuses, and alternatives. *Educational and Psychological Measurement*, 41 (3), 687-699.
- Brunswik, E. (1952). *The conceptual framework of psychology*. Chicago: University of Chicago Press.
- Byrt, T., Bishop, J., & Carlin, J. B. (1993). Bias, prevalence, and kappa. *Journal of Clinical Epidemiology*, 46 (5), 423-429.
- Cicchetti, D. V., & Feinstein, A. R. (1990). High agreement but low kappa: II. Resolving the paradoxes. *Journal of Clinical Epidemiology*, 43(6), 551-558.
- Cohen, J. (1960). A coefficient of agreement for nominal scales, *Educational and Psychological Measurement*, 20(1), 37-46.
- Cohen, J. (1968). Weighted kappa: Nominal scale agreement with provision for scaled disagreement or partial credit. *Psychological Bulletin*, 70, 213-220.
- Crisp, V. (2012). An investigation of rater cognition in the assessment of projects. *Educational Measurement: Issues and Practice*, 31(3), 10-20.
- Engelhard, G. (2013). *Invariant measurement: Using Rasch models in the social, behavioral, and health sciences*. New York: Routledge.
- Engelhard, G., Wang, J., & Wing, S. A. (2018). A tale of two models: Psychometric and cognitive perspectives on rater-mediated assessments using accuracy ratings. *Psychological Test and Assessment Modeling*, 60(1), 33-52.
- Engelhard, G. & Wind, S. (2018). *Invariant measurement with raters and rating scales: Rasch models for rater-mediated assessments*. NY: Routledge.

- Feinstein, A.R. & Cicchetti, D.V. (1990). High agreement but low kappa: I. The problems of two paradoxes. *Journal of Clinical Epidemiology*, 43(6), 543-549.
- Fleiss, J. L. (1981). *Statistical methods for rates and proportions*. NY: Wiley.
- Gwet, K. (2008). Computing inter-rater reliability and its variance in the presence of high Agreement. *British Journal of Mathematical and Statistical Psychology*, DOI: 10.1348/000711006X126600
- Holley, W., & Guilford, J. P. (1964). A note on the G-index of agreement. *Educational and Psychological Measurement*, 24, 749-754.
- Janson, S., & Vegelius, J. (1979). On generalizations of the G index and the phi coefficient to nominal scales. *Multivariate Behavioral Research*, 14, 255-269.
- Joint Committee on the Standards for Educational and Psychological Testing of the American Educational Research Association, the American Psychological Association, and the National Council on Measurement in Education (AERA/APA/NCME, 2014). *The Standards for Educational and Psychological Testing*. Washington, DC: American Psychological Association.
- Krippendorff, K. (1970). Bivariate agreement coefficients for reliability data. In E. R. Borgatta & G. W. Bohrnstedt (Eds.), *Sociological Methodology 1970* (pp. 139-150). San Francisco, CA: Jossey Bass.
- Landis, J. R. & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33, 159-174.
- Linacre, J.M. (1989). *Many-Facets Rasch Measurement*, Chicago, IL: MESA Press.
- Maxwell, A.E. (1977). Coefficients of agreement between observers and their interpretation. *British Journal of Psychiatry*, 130, 79-83.
- Myford, C. M. (2012). Rater cognition research: Some possible directions for the future. *Educational Measurement: Issues and Practice*, 31(3), 48-49.
- Perreault, W. D., & Leigh, L. E. (1989). Reliability of nominal data based on qualitative judgments. *Journal of Marketing Research*, 26, 135-148.
- Porter, J. M. & Jelinek, D. (2011). Evaluating inter-rater reliability of a national assessment model for teacher performance, *International Journal of Educational Policies*, 5(2), 74-87.
- Suto, I. (2012). A critical review of some qualitative research methods used to explore rater cognition. *Educational Measurement: Issues and Practice*, 31(3), 21-30.
<https://doi.org/10.1111/j.1745-399.2012.00240.x>

Thompson, W. D. & Walter, S. D. (1988). A reappraisal of the kappa coefficient. *Journal of Clinical Epidemiology*, 41(10), 949-958.

Xie, Q. (2013). Agree or disagree? A demonstration of an alternative statistic to Cohen's Kappa for measuring the extent and reliability of agreement between observers. Washington, D.C.: Federal Committee on Statistical Methodology, U.S. Office of Management and Budget. Retrieved 5/3/2021 from https://nces.ed.gov/FCSM/pdf/J4_Xie_2013FCSM.pdf.

Table 1

Agreement, leniency, and strictness percentages across both the exact and adjacent agreement conditions

	Exact Agreement			Adjacent Agreement		
	Agreement	Leniency	Strictness	Agreement	Leniency	Strictness
Mean	68.7%	7.7%	23.6%	87.6%	2.4%	10.0%
SD	10.2%	6.5%	12.0%	7.0%	2.8%	7.4%
Minimum	42.0%	0.0%	4.0%	64.0%	0.0%	0.0%
25th percentile	62.0%	3.0%	13.0%	84.0%	0.0%	4.0%
Median	68.0%	6.0%	24.0%	88.0%	2.0%	10.0%
75th percentile	76.0%	12.0%	31.0%	93.0%	5.0%	14.0%
Maximum	88.0%	26.0%	58.0%	100.0%	10.0%	36.0%

Table 2

Chance corrected agreement for the exact agreement condition

	Kappa	Lambda-1	Lambda-2	AC-1	AC-2
Mean	0.486	0.578	0.526	0.606	0.472
SD	0.160	0.140	0.137	0.129	0.092
Minimum	0.121	0.199	0.225	0.273	0.356
25th percentile	0.373	0.487	0.430	0.523	0.359
Median	0.471	0.567	0.519	0.593	0.512
75th percentile	0.592	0.679	0.617	0.704	0.547
Maximum	0.790	0.840	0.797	0.852	0.548

Table 3

Chance corrected agreement for the adjacent agreement condition

	Kappa	Lambda-1	Lambda-2	AC-1	AC-2
Mean	0.666	0.812	0.714	0.848	0.787
SD	0.143	0.107	0.115	0.090	0.052
Minimum	0.315	0.442	0.430	0.532	0.687
25th percentile	0.570	0.757	0.638	0.803	0.763
Median	0.646	0.820	0.689	0.857	0.799
75th percentile	0.768	0.895	0.794	0.918	0.822
Maximum	1.000	1.000	1.000	1.000	0.831

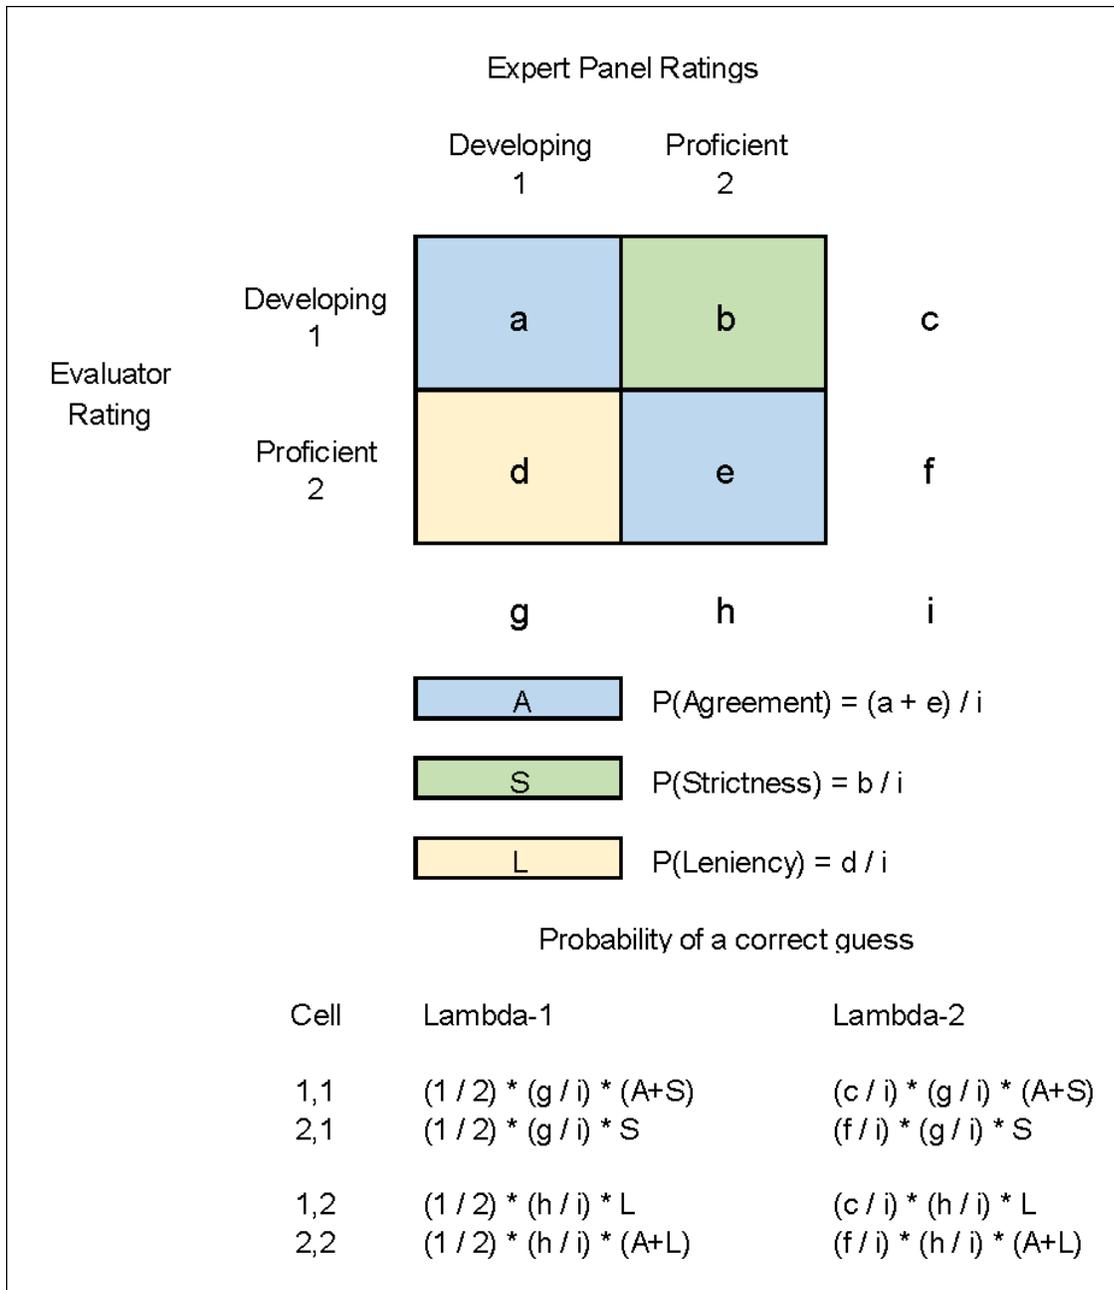


Figure 1. Calculation of Lambda for a 2x2-agreement matrix.

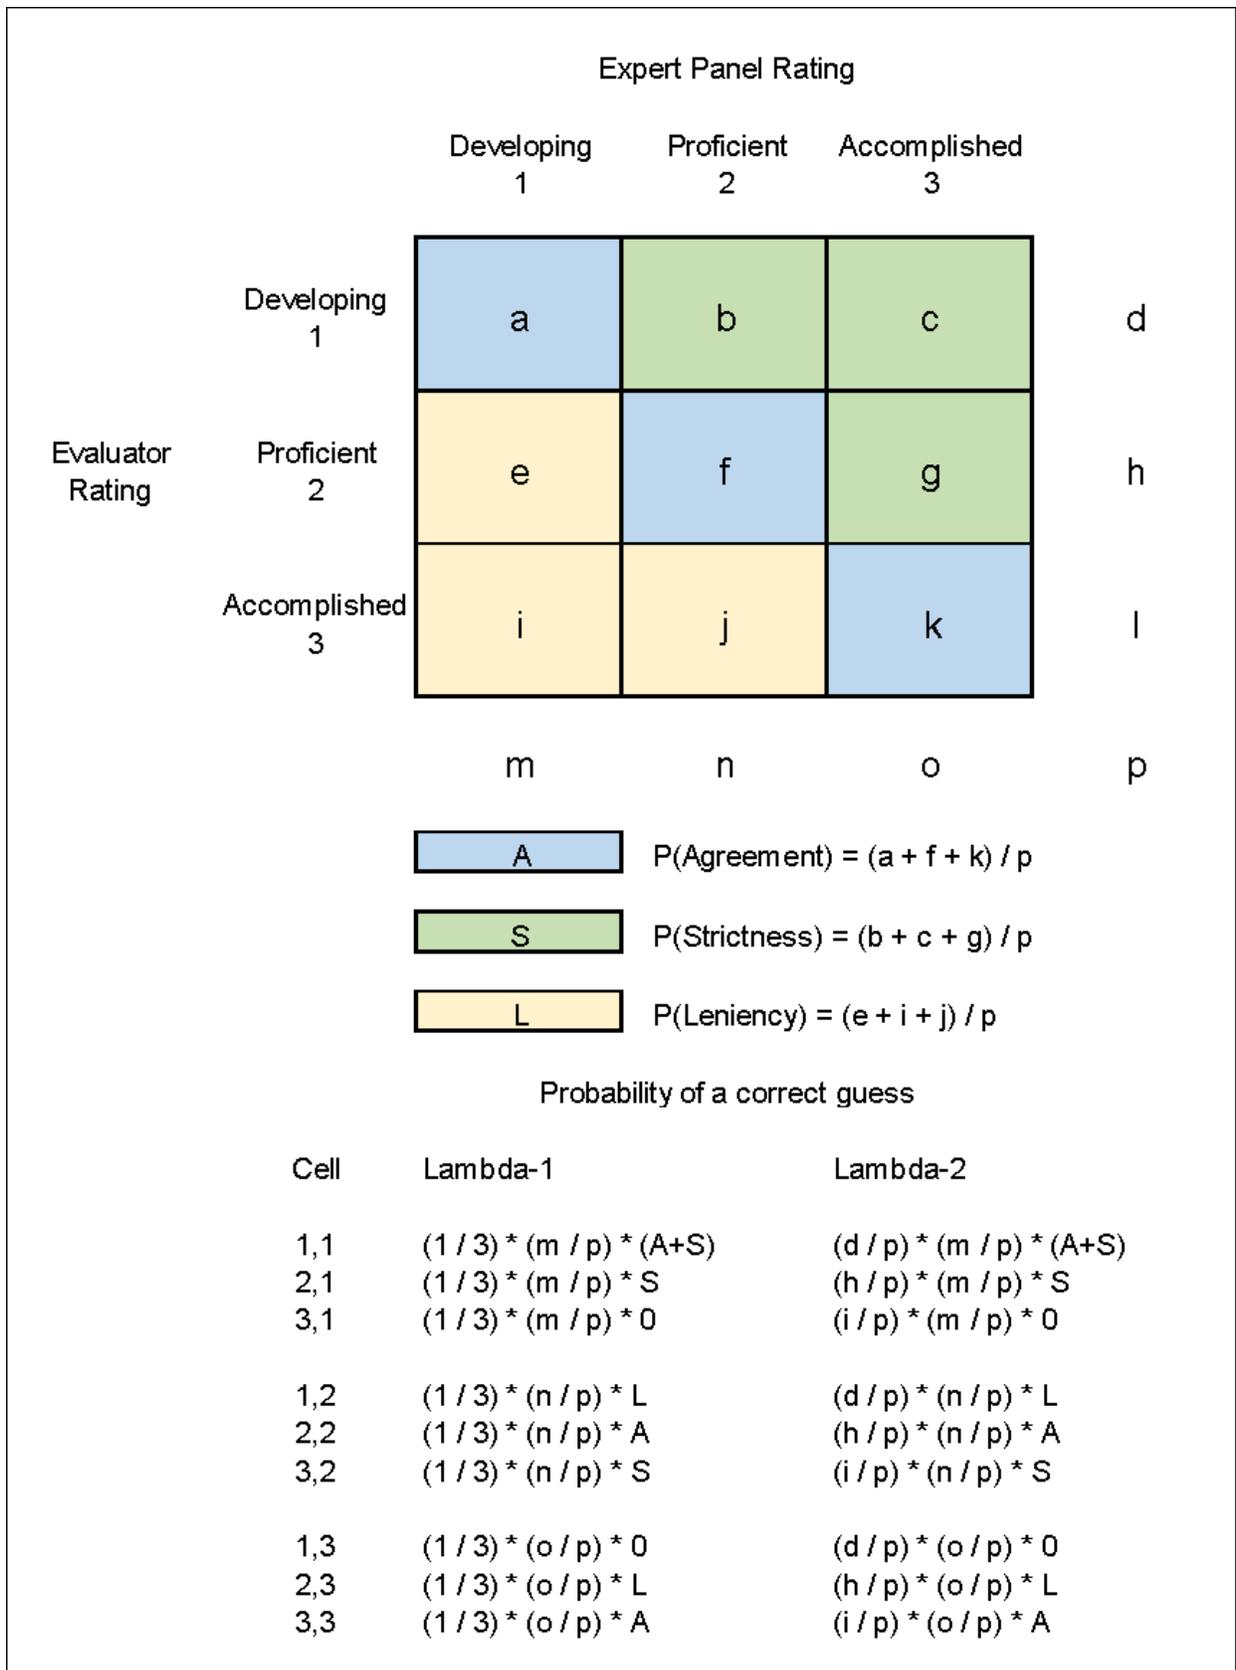


Figure 2. Calculation of Lambda for a 3x3-agreement matrix.

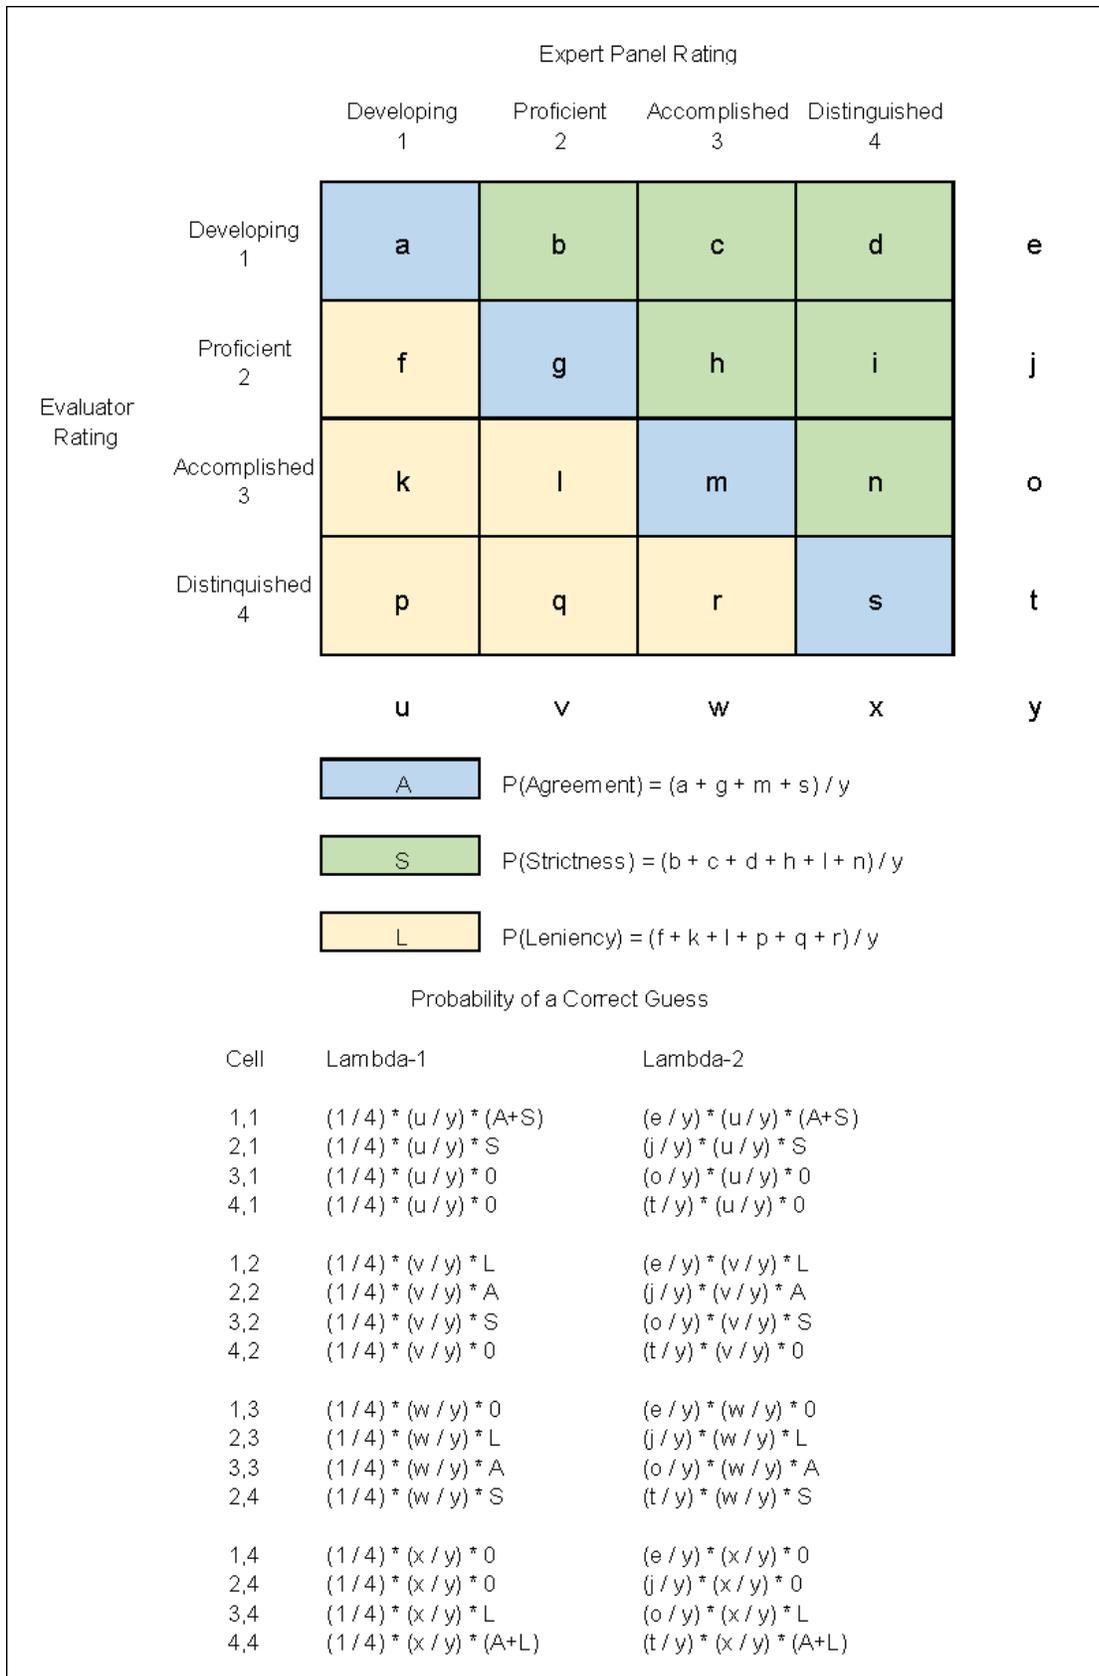


Figure 3. Calculation of Lambda for a 4x4-agreement matrix.

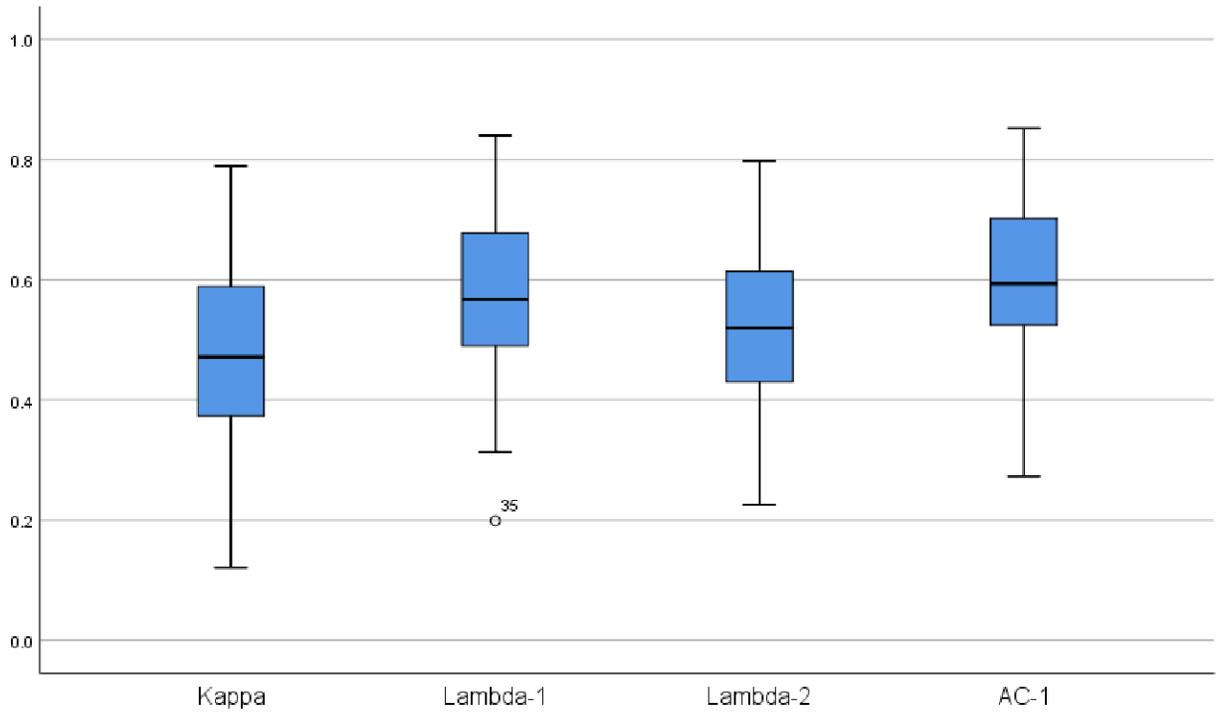


Figure 4. Boxplots of Kappa, Lambda-1, Lambda-2, and AC-1 across all raters in the sample for the exact agreement condition.

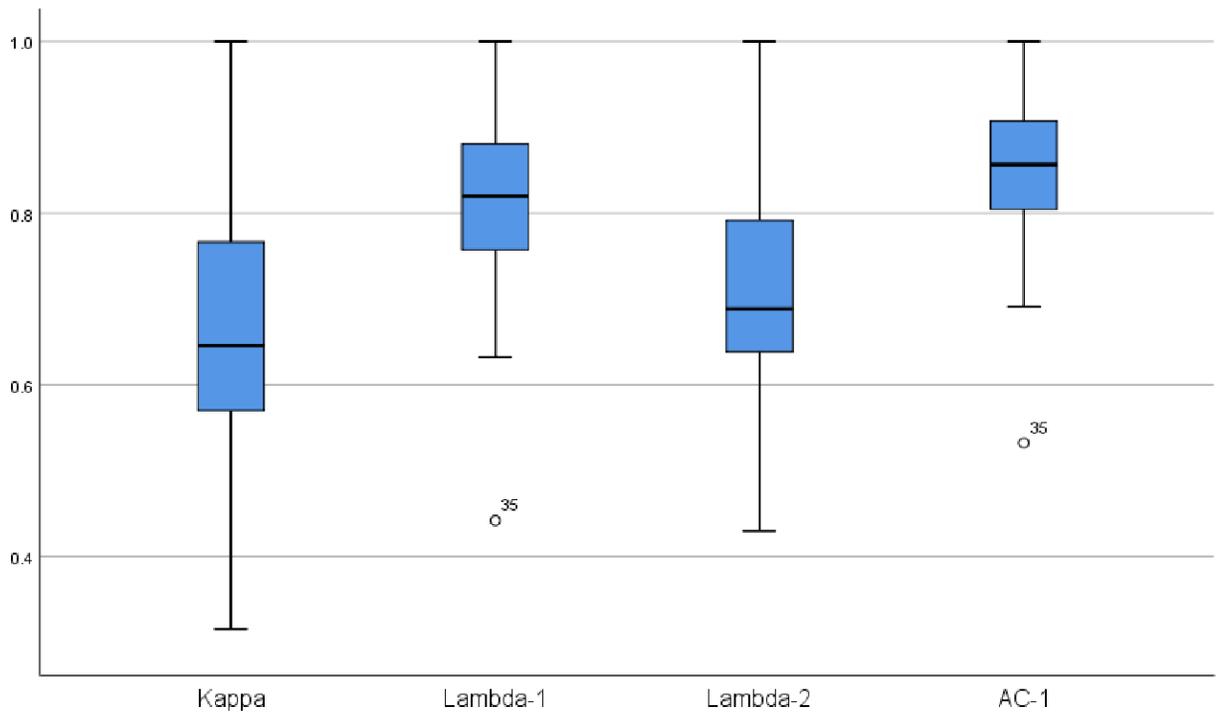


Figure 5. Boxplots of Kappa, Lambda-1, Lambda-2, and AC-1 across all raters in the sample for the adjacent agreement condition.

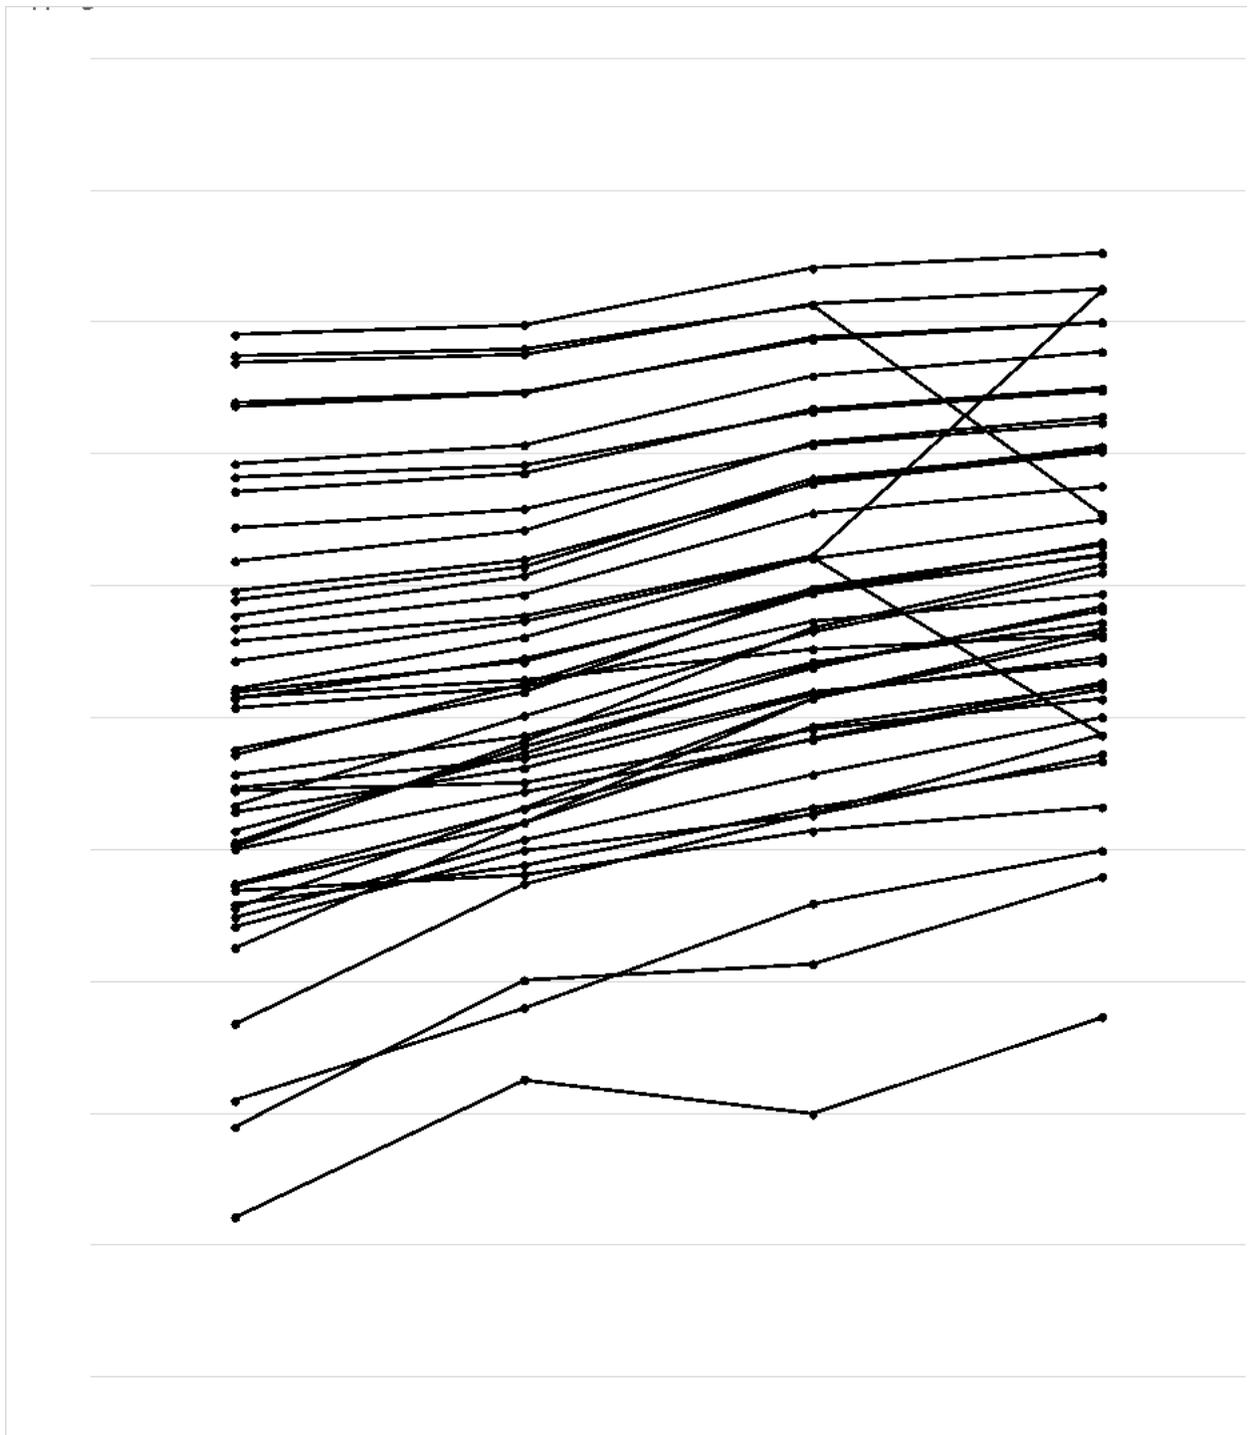


Figure 6. Kappa, Lambda-1, Lambda-2, and AC-1 for each rater under the exact agreement condition.

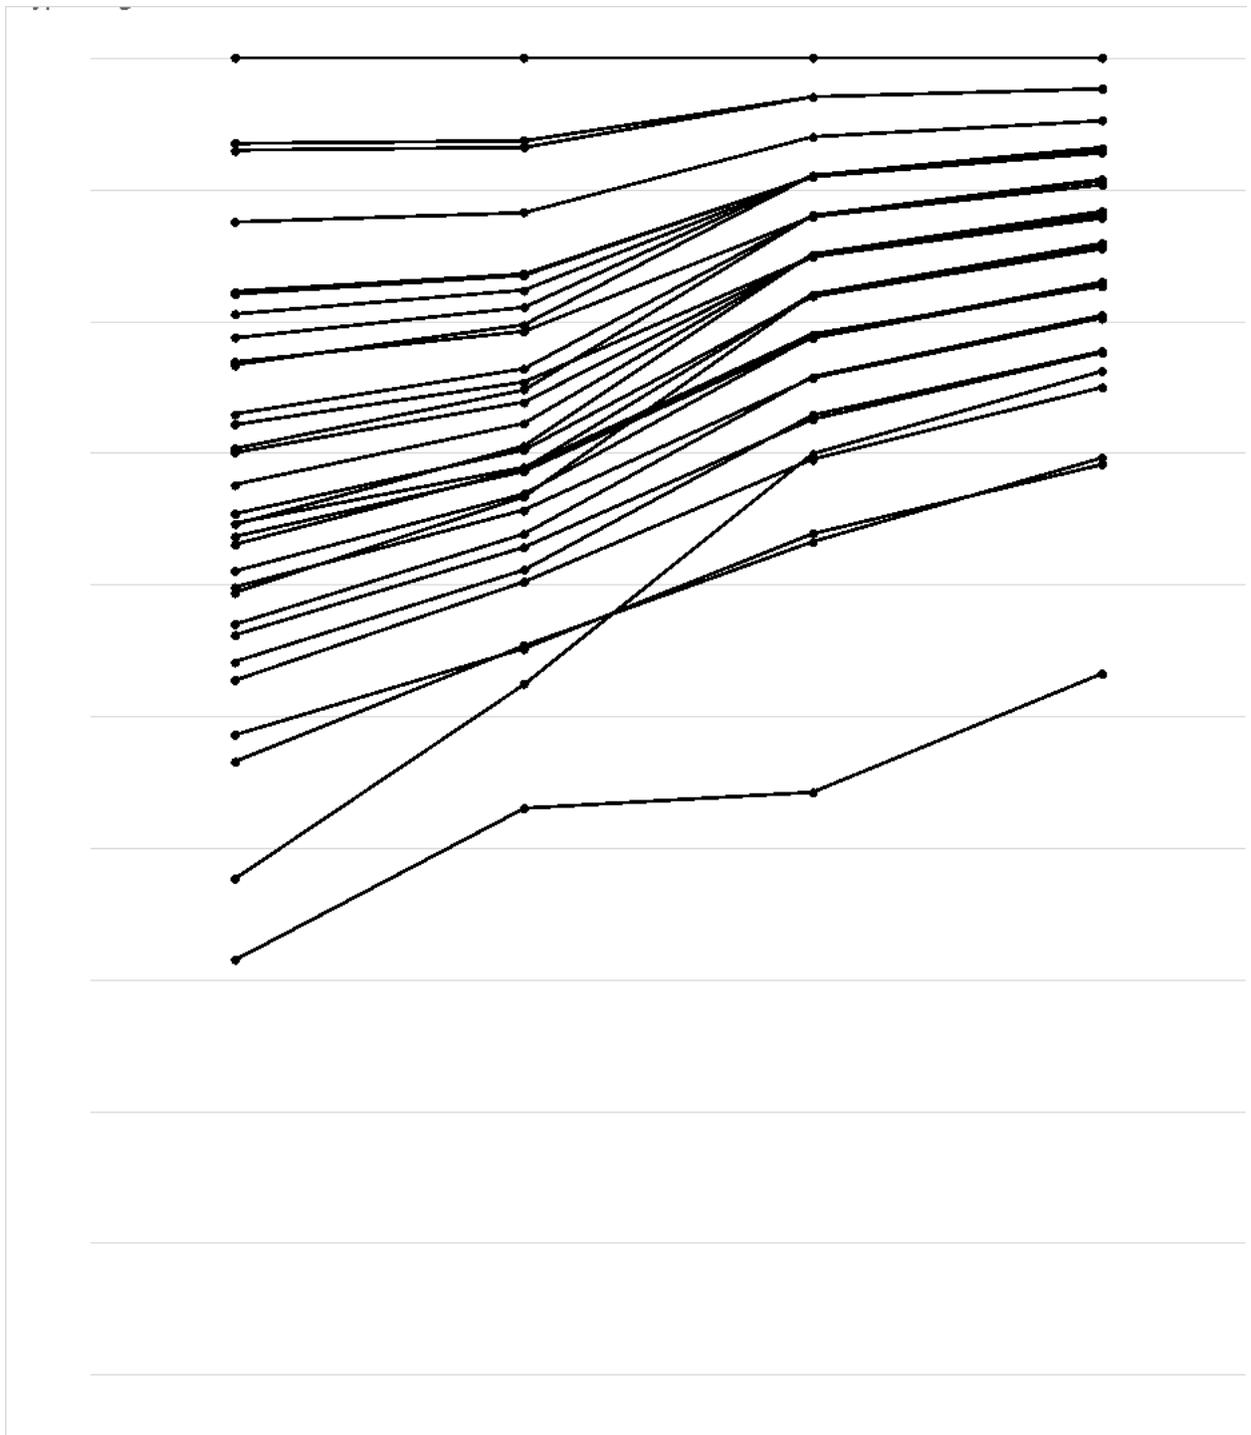


Figure 7. Kappa, Lambda-1, Lambda-2, and AC-1 for each rater under the adjacent agreement condition.

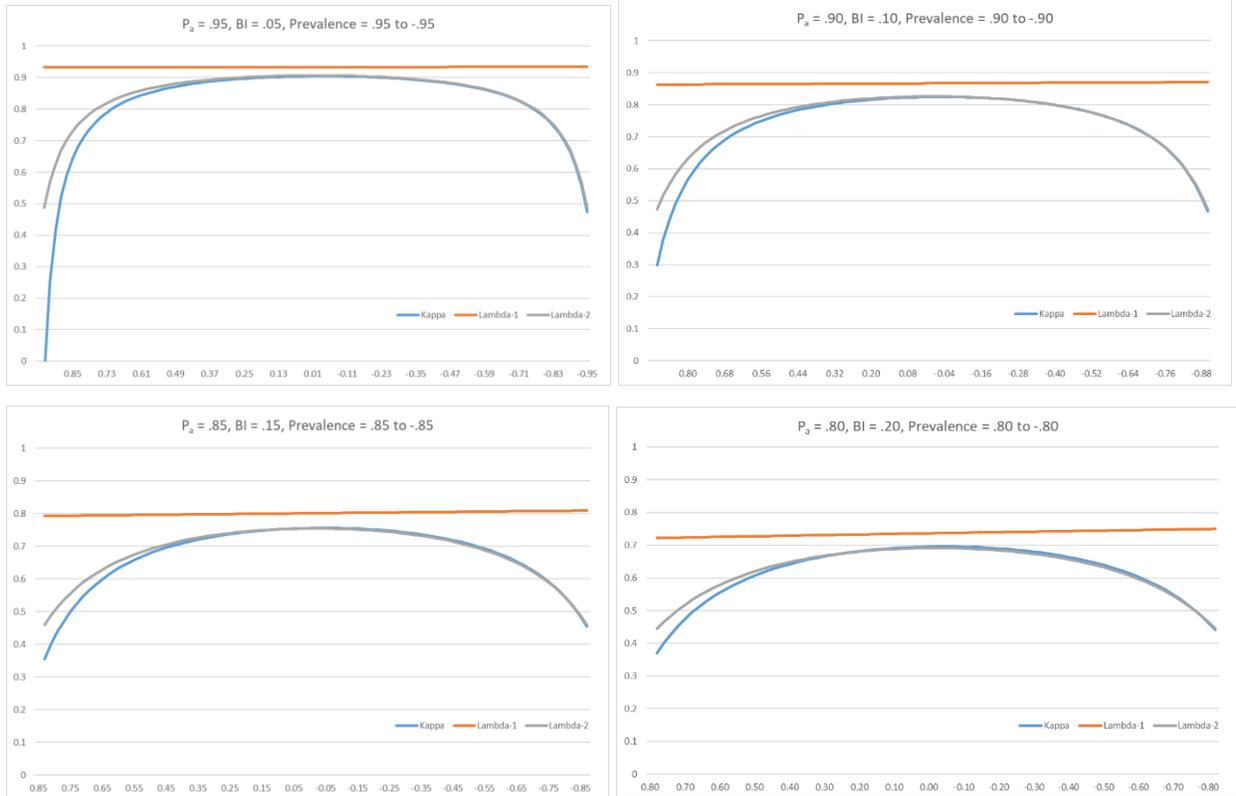


Figure 8. Simulation results. X axis = Prevalence Index, Y axis = Chance-corrected agreement.